

REMARKS

Claims 1 – 22 are pending. Claims 1 – 4 are allowed. Applicant respectfully requests reconsideration of the rejected claims and issuance of a Notice of Allowance.

Applicant thanks Examiner Ellis for the courtesies shown to Applicant's representative during a March 16 personal interview. The substance of the interview is incorporated in the remarks that follow.

On page 2 the Office Action rejects claims 5 – 11 and 13 – 22 under 35 U.S.C. § 103(a) over U.S. Patent 6,081,872 to Matick et al. (hereafter Matick) in view of U.S. Patent 5,713,004 to Kimmel et al. (hereafter Kimmel), Handy, "The Cache Memory Book" (hereafter Handy), and "MC 88410 Secondary cache Controller User's Manual" (hereafter User's Manual). This rejection is respectfully traversed.

Claim 5

The Office Action rejects claim 5 over the combination of Matick, Kimmel, Handy and User's Manual, asserting that these references, when combined, disclose either directly, or through inherency, all the features of claim 5.

Matick is directed to a computer architecture employing a memory hierarchy having multiple cache levels wherein a high level cache (L1) uses SRAM technology and an intermediate level cache (L2) uses DRAM technology without compromising system performance. Because DRAM access is generally much longer than SRAM access, the L2 cache cannot be accessed every processor cycle. Thus, if there is an L1 miss, a L2 directory, which may employ SRAM technology, is accessed on the next processor cycle and an L2 data array waits for the L2 array address on the next (third) processor cycle. Matick discloses several other cache organizations intended to provide good system performance while employing slower DRAM technology. Matick does not disclose or suggest that any of the caches are inclusive, does not disclose snoop processing, and does not, as the Office Action admits, disclose or suggest any other coherency protocol.

The Office Action then asserts that use of inclusive caches is an inherent feature of Matick because for a cache to be non-inclusive, additional steps and/or information must be stored for the caches to know where a piece of data is located in order to maintain coherency. "These additional steps and/or information would be described in the patent if Matick et al. taught a non-inclusive cache design." Applicant strongly disagrees with this logic.

As discussed during the personal interview, in Matick, the address of a piece of data is provided, in the case of the L2 cache, by the L1 directory 11. See Matick, column 4, lines 12 – 26. Thus, Matick directly addresses the supposed “additional steps and/or information” that the Office Action maintains would be needed in a non-inclusive cache design. Furthermore, the Examiner appears to be resorting to speculation as to what Matick would have disclosed if it “taught a non-inclusive cache design.” Applicant respectfully questions how the Examiner could possibly know this. In addition, Matick’s independent claims, including claim 1 for example, appear to be fully enabled by the disclosure, such as at column 4, lines 12 – 26, without resort to any “inherent” design features. Because Matick does not disclose or suggest inclusive intermediate caches, and because such caches are not inherent in Matick’s architecture, Matick can not be used as a reference to anticipate this feature of claim 5.

Moreover, and also as discussed during the personal interview, Matick does not disclose or suggest virtual addressing or the translation of virtual address information into a physical address, and certainly does not disclose or suggest presenting a virtual address for a memory line on a virtually-addressed bus. At most, Matick: (1) merely notes that because a virtual address can appear anywhere in physical memory, such physical memory is defined as fully associative (see column 4, lines 29 – 33); and (2) mentions that the compare circuit 131, which is in the L2 cache directory 13, compares “four separate, large virtual addresses.” The “four separate, large virtual addresses” exist because the L2 cache is four-way set-associative. However, Matick provides no further disclosure about these virtual addresses. In fact, what is clear from Matick’s disclosure is that Matick’s CPU presents, on a point-to-point basis, virtual address information to the L2 cache, and the compare circuit 131 uses the virtual address information to determine if a cache hit has occurred. Note that a “bus” such as recited in Applicant’s claim 5, is not a point-to-point connection. Matick does not disclose explicitly or inherently, putting virtual addresses on a virtually addressed bus for the simple reason that this step is not required for Matick’s system to operate. Furthermore, nowhere does Matick discuss how an address could be translated (for example, by way of a translation lookaside buffer), whether any of the buses to which the cache hierarchy is coupled are in fact virtual buses, or whether virtual address information is presented on a virtual bus, a physical bus, or a hybrid virtual-physical bus. In short, Matick does not disclose or suggest presenting a virtual address for a memory line on a virtually-addressed bus. Kimmel, Handy, and User’s Manual also do not disclose or suggest virtual addressing, and so these references do nothing to cure this defect in Matick.

The Office Action states “[I]t was common in the art to ‘filter out’ snoops at the L2 cache level and not send them to the L1 caches where there was no snoop hit in the L2 cache.” However, the Office Action does not cite any reference for this proposition. In accordance with MPEP 2144.03, the Examiner is respectfully requested to provide such a reference.

Kimmel is directed to a multiprocessor computer architecture that employs snoop processing to reduce the number of invalidate messages on a shared system bus and to maintain cache coherency. The particular snoop protocol is disclosed in User’s Manual, which Kimmel incorporates by reference. Kimmel discloses the use of a hierarchical cache structure, but nowhere does Kimmel disclose or suggest that any of the caches are inclusive. Therefore, Kimmel does not cure Matick’s defects.

Next, the Office Action asserts that Handy teaches inclusive caches that can prescreen invalidation cycles when the data is inclusive. The Office Action concludes that “it would have been obvious to operate the caches of Matick et al. in an inclusive manner for the reasons given by Handy.”

Again, Applicant strongly disagrees with the Office Action’s logic. Matick’s caches employ a specific hardware design. For the L2 cache to be inclusive, it must contain every memory line of higher level caches (i.e., L1 caches). Nowhere does Matick disclose or suggest that the L2 cache has this design feature. Without designed-in inclusivity, Matick’s caches can not simply operate in an inclusive manner. Thus, the Office Action’s assumption concerning the combination of Handy and Matick (and, for that matter, Kimmel) can not be substantiated.

Moreover, use of inclusive caches in Kimmel’s architecture also is not only not inherent, such use would make the disclosed snooping protocol (disclosed in User’s Manual) at least redundant, and probably would mean that Kimmel’s architecture could not maintain cache coherency. In particular, Kimmel relies on comparing global addresses on the shared bus with cache tags. A snoop hit occurs when the global address matches a cache tag for the MTAG (which, in the reference provided with the Office Action is undefined, but presumably is a secondary level cache). A primary cache snoop hit occurs when the cache tag for a valid primary data cache matches the global address on the shared bus. This protocol is incompatible with the claimed snoop processing and inclusive intermediate cache design, recited in claim 5.

In contrast to Matick, Kimmel, Handy, and User’s Manual, claim 5 recites:

“establishing the intermediary caches as inclusive
caches . . . ;
presenting a virtual address for a memory line on a
virtually-addressed bus;
initiating snoop processing of the intermediary caches
. . .”

As discussed above, Matick, does not disclose or suggest intermediary caches as inclusive caches, nor are inclusive intermediary caches inherent in Matick. Matick also does not disclose or suggest presenting a virtual address for a memory line on a virtually-addressed bus. In the invention recited in claim 5, the CPU initiates the snoop processing (“initiating snoop processing of the intermediary caches”) but needs the virtual address information to do so. Matick does not disclose or suggest snoop processing, and hence does not require the same process as recited in claim 5. Kimmel, Handy, and User’s Manual do nothing to cure these defects in Matick.

Thus, for all the reasons stated above, Matick, Kimmel, Handy and User’s Manual, individually and in combination, do not disclose or suggest all the features of claim 5, and claim 5 is patentable.

Claims 9 and 19

Claims 9 and 19 each contain means-plus-function elements. Analyzing these claims requires reference to the specification to determine the disclosed structure. In claim 9, a means-plus-function element recites “means for ensuring the intermediary cache is inclusive of associated local caches.” Claim 19 recites “means for reducing processing time and bus bandwidth during snoop processing.” The specification discloses at least three distinct “structures” to perform the claimed functions. First, the capacity of the intermediate level cache is always at least that of the total capacity of all the associated upper level caches. Second, a snoop protocol is used such that if a line is evicted from the intermediate level cache, the line is also evicted from any upper level cache that also holds that line. Note that if a snoop hit occurs with the intermediate level cache, a further snoop operation is required to determine which upper level cache holds that line. Third, the intermediate level caches are implemented as coherency filters, which are simply data-less caches. These coherency filters store virtual address bits, but not the data or instructions related to the virtual addresses. See page 10, lines 11 – 30.

As discussed above with respect to claim 5, Matick, Kimmel, Handy, and User’s Manual, individually and in combination, do not disclose or suggest the structures that

perform the functions recited in claims 9 and 19, and therefore claims 9 and 19 are patentable.

Claims 6 – 8, 10, 11, 13 – 18, and 20 – 22

Claims 6 – 8 depend from patentable claim 5, claims 10, 11, and 13 – 18 depend from patentable claim 9, and claims 20 – 22 depend from patentable claim 19. For these reasons and the additional features they recite, claims 6 – 8, 10, 11, 13 – 18, and 20 – 22 are also patentable. Withdrawal of the rejection of claims 5 – 11 and 13 – 22 under 35 U.S.C. § 103(a) is respectfully requested.

Claim 12

On page 5 the Office Action rejects claim 12 under 35 U.S.C. § 103(a) over Matick in view of Kimmel, Handy, User's Manual, and further in view of U.S. Patent 5,778,424 to Guy. This rejection is respectfully traversed.

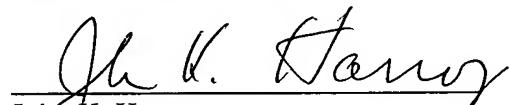
Claim 12 depends from patentable claim 9. For this reason and the additional features it recites, claim 12 is also patentable. Withdrawal of the rejection of claim 12 under 35 U.S.C. § 103(a) is respectfully requested.

In view of the above remarks, Applicant respectfully submits that the application is in condition for allowance. Prompt examination and allowance are respectfully requested.

Should the Examiner believe that anything further is desired in order to place the application in even better condition for allowance, the Examiner is invited to contact Applicant's undersigned representative at the telephone number listed below.

Respectfully submitted,

Date: March 16, 2005


John K. Harrop
Registration No. 41,817
Andrews Kurth LLP
1701 Pennsylvania Ave, N.W.
Suite 300
Washington, DC 20006
Tel. (202) 662-3050
Fax (202) 662-2739